

US Patent 5,611,049, Claim 1:

1. In a network of digital computers that includes a first plurality of Network Distributed Cache ("NDC") sites, each NDC site including an NDC that has an NDC buffer, a method for projecting images of a stored dataset from an NDC server terminator site into a second plurality of NDC client terminator sites in response to requests to concurrently access such stored dataset transmitted from a third plurality of client sites respectively to the second plurality of NDC client terminator sites, the method comprising the steps of:

(a) the NDC receiving the request to access data in the stored dataset;

(b) the NDC checking the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there;


(c) if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if the NDC site receiving the request is not the NDC server terminator site for the stored dataset, the NDC of this NDC site transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site;

(d) if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if the NDC site receiving the request is the NDC server terminator site for the stored dataset, the NDC of the NDC server terminator site accessing the stored dataset to project an image of the requested data into the NDC buffer of the NDC server terminator site;

(e) repeating the steps (a) through (d) until the NDC buffer of the downstream NDC site receiving the request contains a projected image of all requested data;

(f) each successive NDC site, having obtained a projected image of all the requested data, returning the requested data upstream to the NDC site from which the NDC site received the request until the requested data arrives at the NDC client terminator site, each NDC site that returns data upstream to the requesting NDC site retaining a copy of the returned data that the returning NDC site may subsequently transmit to an NDC site other than the NDC site to which the returning NDC site first returned the data, whereby images of the stored dataset may be projected concurrently from a single NDC site into the second plurality of NDC client terminator sites; and

(g) the NDC client terminator site, upon receiving the requested data, returning the requested data to the client site that requested access to the stored dataset.




How Proxy Cache Works

Figure 5 illustrates how BorderManager caches HTML documents and other cacheable content.

Figure 5: A proxy cache saves repeatedly-used objects to speed access and reduce Internet traffic.

- (a) the NDC receiving the request to access data in the stored dataset;
- (b) the NDC checking the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there;




1. A browser issues a request for a file named DOC.HTML. This request is sent to the proxy cache over a 10 Mbps Ethernet LAN segment. In this case, the request results in a "cache miss" because the proxy cache has never serviced a request for that document before.
2. The proxy cache initiates a request for DOC.HTML from the origin web server on behalf of the browser. This request is sent over a T1 line to an ISP, then traverses the Internet until it arrives at the origin server.
3. The origin web server responds to the proxy's request by sending DOC.HTML. This transmission is much faster than a response to a browser due to the proxy's optimized receive window that can receive up to 64KB at one time and stays open to receive multiple responses. The proxy then places DOC.HTML in its cache.
4. The proxy cache responds to the original browser request with DOC.HTML.
5. Now when the same browser (or any other browser) issues a request for DOC.HTML, the request results in a "cache hit" because the proxy has kept a copy of the document in its cache.
6. In this case, the proxy replies immediately to the browser request because it has DOC.HTML in cache. The proxy's response is transmitted at 10 Mbps to the browser, eliminating the need to fetch the document again from the origin server on the Internet.

How Proxy Cache Works

Figure 5 illustrates how BorderManager caches HTML documents and other cacheable content.

- (a) the NDC receiving the request to access data in the stored dataset;
- (b) the NDC checking the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there;

Figure 5: A proxy cache saves repeatedly-used objects to speed access and reduce Internet traffic.



1. A browser issues a request for a file named DOC.HTML. This request is sent to the proxy cache over a 10 Mbps Ethernet LAN segment. In this case, the request results in a "cache miss" because the proxy cache has never serviced a request for that document before.
2. The proxy cache initiates a request for DOC.HTML from the origin web server on behalf of the browser. This request is sent over a T1 line to an ISP, then traverses the Internet until it arrives at the origin server.
3. The origin web server responds to the proxy's request by sending DOC.HTML. This transmission is much faster than a response to a browser due to the proxy's optimized receive window that can receive up to 64KB at one time and stays open to receive multiple responses. The proxy then places DOC.HTML in its cache.
4. The proxy cache responds to the original browser request with DOC.HTML.
5. Now when the same browser (or any other browser) issues a request for DOC.HTML, the request results in a "cache hit" because the proxy has kept a copy of the document in its cache.
6. In this case, the proxy replies immediately to the browser request because it has DOC.HTML in cache. The proxy's response is transmitted at 10 Mbps to the browser, eliminating the need to fetch the document again from the origin server on the Internet.